

Parallelization of continuous and discontinuous Galerkin dual-primal Isogeometric tearing and interconnecting methods

Christoph Hofer¹

¹ Johannes Kepler University (JKU), Altenbergerstr. 69, A-4040 Linz, Austria,
christoph.hofer@jku.at

Abstract. In this paper we investigate the parallelization of dual-primal isogeometric tearing and interconnecting (IETI-DP) type methods for solving large-scale continuous and discontinuous Galerkin systems of equations arising from Isogeometric analysis of elliptic boundary value problems. These methods are extensions of the finite element tearing and interconnecting methods to isogeometric analysis. The algorithms are implemented by means of energy minimizing primal subspaces. We discuss how these methods can efficiently be parallelized in a distributed memory setting. Weak and strong scaling studies presented for two and three dimensional problems show an excellent parallel efficiency.

Key words: Diffusion problems, Isogeometric analysis, discontinuous Galerkin, IETI-DP, parallelization, MPI

1 Introduction

Isogeometric Analysis (IgA) is a novel methodology for the numerical solution of partial differential equations (PDE). IgA was first introduced by Hughes, Cottrell and Bazilevs in [25], see also the monograph [7] for a comprehensive presentation of the IgA framework and the recent survey article [3]. The main principle is to use the same basis functions for describing the geometry and to represent the discrete solution of the PDE problem under consideration. The most common choices are B-Splines, Non Uniform Rational B-Splines (NURBS), T-Splines, Truncated Hierarchical B-Splines (THB-Splines), etc., see, e.g., [15], [16] and [2]. One of the strengths of IgA is the capability of creating high-order splines spaces, while keeping the number of degrees of freedom quite small. Moreover, having basis functions with high smoothness is useful when considering higher-order PDEs, e.g., the biharmonic equation.

In many cases the domain can not be represented with a single mapping, referred to as *geometrical mapping*. Complicated geometries are decomposed into simple domains, called *patches*, which are topologically equivalent to a cube. The set of patches forming the computational domain is called multipatch domain. The obtained patch parametrizations and the original geometry may not be identical. The result are small gaps and overlaps occurring at the interfaces of the patches, called *segmentation crimes*, see [26], [42] and [24] for a comprehensive analysis. Nevertheless, one still wants to solve PDEs on such domains. To do so, numerical schemes based on the discontinuous Galerkin (dG) method for elliptic PDEs were developed in [21], [23] and [22]. There, the corresponding error analysis is also provided. In addition to domains with segmentation crimes, the dG formulation is very useful when considering different B-Splines spaces on each patch, e.g., non-matching grids at the interface and different spline degrees. An analysis for the dG-IgA formulation with

extensions to low regularity solutions can be found in [38]. For a detailed discussion of dG for finite element methods, we refer, e.g., to [45] and [8].

In the present paper, we are considering fast solution methods for linear systems arising from the discretization of elliptic PDEs by means of IgA. We investigate non-overlapping domain decomposition (DD) methods of the dual-primal tearing and interconnecting type. This type of methods are equivalent to the so called Balancing Domain Decomposition by Constraints (BDDC) methods, see [40], [46], [43] and references therein. The version based on a conforming Galerkin (cG) discretization, called dual-primal isogeometric tearing and interconnecting (IETI-DP) method was first introduced in [35] and the equivalent IgA BDDC method was analysed in [5]. Further extensions to the analysis are presented in [19]. The version based on the dG formulation, abbreviated by dG-IETI-DP, was introduced in [20] and analyzed in [18], see [12], [13] and [14] for the corresponding finite element counterparts. We also want to mention development in overlapping Schwarz methods, see, e.g., [4] and [6]. The aim of this paper is to present the parallel scalability of the cG and dG IETI-DP methods. We investigate weak and strong scaling in two and three dimensional domains for different B-Spline degrees. The implemented algorithms are based on energy minimizing primal subspaces, which simplifies the parallelization of the solver part, but having more effort in the setup phase (assembling phase). We rephrase key parts of this algorithm and discuss how to realize the communication by means of Message Passing Interface (MPI). In general, FETI-DP and equivalent BDDC methods are by nature well suited for large-scale parallelization and has been widely studied for solving large-scale finite element equations, e.g., in [33], [44], [31] and [28], see also [30] for a hybrid OpenMP/MPI version. Considering a domain decomposition with several ten thousands of subdomains, the influence of the coarse grid problem becomes more and more significant. Especially, its LU-factorization is the bottleneck of the algorithm. The remedy is to reformulate the FETI-DP system in such a way that the solution of the coarse grid problem is not required in the application of the system matrix, but in the preconditioner. This enables the use of inexact methods like geometric or algebraic multigrid, see, e.g., [29], [28], [32], [33] and [34]. Moreover, inexact solvers can also be used in the scaled Dirichlet preconditioner and, if using the saddle point formulation, also for the local solvers, cf., [32], see also [33], [44] and references therein for alternative approaches by means of hybrid FETI. We also want to mention inexact version for the BDDC method, see, e.g., [47], [48], [10], [39] and [49]. FETI-DP methods has also been successfully applied to non-linear problems by means of a non-linear version of FETI-DP. We want to highlight recent advances presented, e.g., in [27], [29] and [28], showing excellent scalability on large-scale supercomputers.

In the present paper, we consider the following second-order elliptic boundary value problem in a bounded Lipschitz domain $\Omega \subset \mathbb{R}^d$, with $d \in \{2, 3\}$: Find $u : \overline{\Omega} \rightarrow \mathbb{R}$ such that

$$-\operatorname{div}(\alpha \nabla u) = f \text{ in } \Omega, \quad u = 0 \text{ on } \Gamma_D, \quad \text{and } \alpha \frac{\partial u}{\partial n} = g_N \text{ on } \Gamma_N, \quad (1)$$

with given, sufficient smooth data f, g_N and α , where the coefficient α is uniformly bounded from below and above by some positive constants α_{\min} and α_{\max} , respectively. The boundary $\partial\Omega$ of the computational domain Ω consists of a Dirichlet part Γ_D of positive boundary measure and a Neumann part Γ_N . Furthermore, we assume that the Dirichlet boundary Γ_D is always a union of complete patch sides (edges / face in 2D / 3D) which are uniquely defined in IgA. Without loss of generality, we assume homogeneous Dirichlet conditions. This can always be obtained by homogenization. By means of inte-

gration by parts, we arrive at the weak formulation of (1) which reads as follows: Find $u \in V_D = \{u \in H^1 : \gamma_0 u = 0 \text{ on } \Gamma_D\}$ such that

$$a(u, v) = \langle F, v \rangle \quad \forall v \in V_D, \quad (2)$$

where γ_0 denotes the trace operator. The bilinear form $a(\cdot, \cdot) : V_D \times V_D \rightarrow \mathbb{R}$ and the linear form $\langle F, \cdot \rangle : V_D \rightarrow \mathbb{R}$ are given by the expressions

$$a(u, v) := \int_{\Omega} \alpha \nabla u \cdot \nabla v \, dx \quad \text{and} \quad \langle F, v \rangle := \int_{\Omega} f v \, dx + \int_{\Gamma_N} g_N v \, ds.$$

The remainder of the paper is organized as follows: In Section 2, we give a short introduction to isogeometric analysis, providing the basic definitions and notations. Section 3 describes the different discretizations of the model problem obtained the continuous and discontinuous Galerkin methods. In Section 4, we formulate the IETI-DP method for both discretizations and provide implementation details. The way how the algorithm is parallelized is explained in Section 5. Numerical examples are presented in Section 6. Finally we draw some conclusions in Section 7.

2 Isogeometric Analysis

In this section, we give a very short overview about IgA. For a more comprehensive study, we refer to, e.g., [7] and [38].

Let $\hat{\Omega} := (0, 1)^d$, $d \in \{2, 3\}$, be the d -dimensional unit cube, which we refer to as the *parameter domain*. Let p_ι and M_ι , $\iota \in \{1, \dots, d\}$, be the B-Spline degree and the number of basis functions in x_ι -direction. Moreover, let $\Xi_\iota = \{\xi_1 = 0, \xi_2, \dots, \xi_{n_\iota} = 1\}$, $n_\iota = M_\iota - p_\iota - 1$, be a partition of $[0, 1]$, called *knot vector*. With this ingredients we are able to define the B-Spline basis $\hat{N}_{i,p}$, $i \in \{1, \dots, M_\iota\}$ on $[0, 1]$ via Cox-De Boor's algorithm, cf. [7]. The generalization to $\hat{\Omega}$ is realized by considering a tensor product, again denoted by $\hat{N}_{i,p}$, where $i = (i_1, \dots, i_d)$ and $p = (p_1, \dots, p_d)$ are a multi-indices. For notational simplicity, we define $\mathcal{I} := \{(i_1, \dots, i_d) \mid i_\iota \in \{1, \dots, M_\iota\}\}$ as the set of multi-indices. Since the tensor product knot vector Ξ provides a partition of $\hat{\Omega}$, it introduces a mesh $\hat{\mathcal{Q}}$, and we denote a mesh element by \hat{Q} , called *cell*.

The B-Spline basis functions parametrize the computational domain Ω , also called *physical domain*. It is given as image of parameter domain under the *geometrical mapping* $G : \hat{\Omega} \rightarrow \mathbb{R}^d$, defined as

$$G(\xi) := \sum_{i \in \mathcal{I}} P_i \hat{N}_{i,p}(\xi),$$

with the control points $P_i \in \mathbb{R}^d$, $i \in \mathcal{I}$. The image of the mesh $\hat{\mathcal{Q}}_h$ under G defines the mesh on Ω , denoted by \mathcal{Q}_h with cells Q . Both meshes possess a characteristic mesh size \hat{h} and h , respectively. More complicated geometries Ω have to be represented with multiple non-overlapping domains $\Omega^{(k)} := G^{(k)}(\hat{\Omega})$, $k = 1, \dots, N$, called *patches*, where each patch is associated with a different geometrical mapping $G^{(k)}$. We sometimes call $\overline{\Omega} := \bigcup_{k=1}^N \overline{\Omega}^{(k)}$ a *multipatch domain*. Furthermore, we denote the set of all indices l such that $\Omega^{(k)}$ and $\Omega^{(l)}$ have a common interface $F^{(kl)}$ by $\mathcal{I}_{\mathcal{F}}^{(k)}$. We define the interface $\Gamma^{(k)}$ of $\Omega^{(k)}$ as $\Gamma^{(k)} := \bigcup_{l \in \mathcal{I}_{\mathcal{F}}^{(k)}} F^{(kl)}$.

We use B-Splines not only for defining the geometry, but also for representing the approximate solution of our PDE. This motivates to define the basis functions in the physical space $N_{i,p} := \hat{N}_{i,p} \circ G^{-1}$ and the corresponding discrete space as

$$V_h := \text{span}\{N_{i,p}\}_{i \in \mathcal{I}}. \quad (3)$$

Moreover, each function $u_h(x) = \sum_{i \in \mathcal{I}} u_i N_{i,p}(x)$ is associated with the coefficient vector $\mathbf{u} = (u_i)_{i \in \mathcal{I}}$. This map is known as *Ritz isomorphism* or *IgA isomorphism* in connection with IgA. One usually writes this relation as $u_h \leftrightarrow \mathbf{u}$. In the following, we will use the notation u_h for the function and its vector representations. If we consider a single patch $\Omega^{(k)}$ of a multipatch domain Ω , we will use the notation $V_h^{(k)}, N_{i,p}^{(k)}, \hat{N}_{i,p}^{(k)}, G^{(k)}, \dots$ with the analogous definitions. To keep notation simple, we will use h_k and \hat{h}_k instead of $h^{(k)}$ and $\hat{h}^{(k)}$, respectively.

3 Galerkin Methods for Isogeometric Analysis

In this section we rephrase the variational formulation for the continuous and discontinuous Galerkin method for multipatch IgA systems.

3.1 Continuous Galerkin method

We are considering the finite dimensional subspace V_h^{cG} of V_D , where V_h^{cG} is given by

$$V_h^{cG} := \{v \mid v|_{\Omega^{(k)}} \in V_h^{(k)}\} \cap H^1(\Omega).$$

Since, we restrict ourselves to homogeneous Dirichlet conditions, we look for the Galerkin approximate u_h from $V_{D,h}^{cG} \subset V_h^{cG}$, where $V_{D,h}^{cG}$ contains all functions, which vanish on the Dirichlet boundary. The Galerkin IgA scheme reads as follows: Find $u_h \in V_{D,h}^{cG}$ such that

$$a(u_h, v_h) = \langle F, v_h \rangle \quad \forall v_h \in V_{D,h}. \quad (4)$$

There exists a unique IgA solution $u_h \in V_{D,h}^{cG}$ of (4) that converges to the solution $u \in V_D$ of (2) if h tends to 0. Due to Cea's lemma, the usual discretization error estimates in the H^1 - norm follow from the corresponding approximation error estimates, see [1] or [3].

3.2 Discontinuous Galerkin method

In the dG-IgA scheme, we again use the spaces $V_h^{(k)}$ of B-Splines defined in (3), whereas now discontinuities are allowed across the patch interfaces $F^{(kl)}$. The continuity of the function value and its normal fluxes are then enforced in a weak sense by adding additional terms to the bilinear form. We define the dG-IgA space

$$V_h^{dG} := \{v \mid v|_{\Omega^{(k)}} \in V_h^{(k)}\}, \quad (5)$$

where $V_h^{(k)}$ is defined as in (3). A comprehensive study of dG schemes for FE can be found in [45] and [8]. For an analysis of the dG-IgA scheme, we refer to [38]. We define $V_{D,h}^{dG}$ as the space of all functions from V_h that vanish on the Dirichlet boundary Γ_D . Having these

definitions at hand, we can define the discrete problem based on the Symmetric Interior Penalty (SIP) dG formulation as follows: Find $u_h \in V_{D,h}^{dG}$ such that

$$a_h(u_h, v_h) = \langle F, v_h \rangle \quad \forall v_h \in V_{D,h}^{dG}, \quad (6)$$

where

$$a_h(u, v) := \sum_{k=1}^N a_e^{(k)}(u, v) \quad \text{and} \quad \langle F, v \rangle := \sum_{k=1}^N \left(\int_{\Omega^{(k)}} f v^{(k)} dx + \int_{\Gamma_N^{(k)}} g_N v^{(k)} ds \right),$$

$$a_e^{(k)}(u, v) := a^{(k)}(u, v) + s^{(k)}(u, v) + p^{(k)}(u, v),$$

and

$$a^{(k)}(u, v) := \int_{\Omega^{(k)}} \alpha^{(k)} \nabla u^{(k)} \nabla v^{(k)} dx,$$

$$s^{(k)}(u, v) := \sum_{l \in \mathcal{I}_{\mathcal{F}}^{(k)}} \int_{F^{(kl)}} \frac{\alpha^{(k)}}{2} \left(\frac{\partial u^{(k)}}{\partial n} (v^{(l)} - v^{(k)}) + \frac{\partial v^{(k)}}{\partial n} (u^{(l)} - u^{(k)}) \right) ds,$$

$$p^{(k)}(u, v) := \sum_{l \in \mathcal{I}_{\mathcal{F}}^{(k)}} \int_{F^{(kl)}} \frac{\delta \alpha^{(k)}}{h_{kl}} (u^{(l)} - u^{(k)}) (v^{(l)} - v^{(k)}) ds.$$

Here the notation $\frac{\partial}{\partial n}$ denotes the derivative in the direction of the outer normal vector, δ a positive sufficiently large penalty parameter, and h_{kl} the harmonic average of the adjacent mesh sizes, i.e., $h_{kl} = 2h_k h_l / (h_k + h_l)$.

We equip $V_{D,h}^{dG}$ with the dG-norm

$$\|u\|_{dG}^2 = \sum_{k=1}^N \left[\alpha^{(k)} \|\nabla u^{(k)}\|_{L^2(\Omega^{(k)})}^2 + \sum_{l \in \mathcal{I}_{\mathcal{F}}^{(k)}} \frac{\delta \alpha^{(k)}}{h_{kl}} \int_{F^{(kl)}} (u^{(k)} - u^{(l)})^2 ds \right]. \quad (7)$$

Furthermore, we define the bilinear forms

$$d_h(u, v) = \sum_{k=1}^N d^{(k)}(u, v) \quad \text{where} \quad d^{(k)}(u, v) = a^{(k)}(u, v) + p^{(k)}(u, v),$$

for later use. We note that $\|u_h\|_{dG}^2 = d_h(u_h, u_h)$.

Lemma 1. *Let δ be sufficiently large. Then there exist two positive constants γ_0 and γ_1 , which are independent of $h_k, H_k, \delta, \alpha^{(k)}$ and u_h such that the inequalities*

$$\gamma_0 d^{(k)}(u_h, u_h) \leq a_e^{(k)}(u_h, u_h) \leq \gamma_1 d^{(k)}(u_h, u_h), \quad \forall u_h \in V_{D,h}^{dG} \quad (8)$$

are valid for all $k = 1, 2, \dots, N$. Furthermore, we have the inequalities

$$\gamma_0 \|u_h\|_{dG}^2 \leq a_h(u_h, u_h) \leq \gamma_1 \|u_h\|_{dG}^2, \quad \forall u_h \in V_{D,h}^{dG}. \quad (9)$$

This Lemma is an equivalent statement of Lemma 2.1 in [13] for IgA, and the proof can be found in [20]. A direct implication of (9) is the well posedness of the discrete problem (6) by the Theorem of Lax-Milgram. The consistency of the method together with the interpolation estimates of B-Splines lead to the a-priori error estimate, established in [38]. We note that, in [38], the results were obtained for the Incomplete Interior Penalty (IIP) scheme. An extension to SIP-dG and the use of harmonic averages for h and/or α are discussed in Remark 3.1 in [38], see also [37].

For both the cG and dG formulation, we choose the B-Spline function $\{N_{i,p}\}_{i \in \mathcal{I}_0}$ as basis for the space V_h^X , $X \in \{cG, dG\}$, where \mathcal{I}_0 contains all indices of \mathcal{I} , where the corresponding basis functions do not have a support on the Dirichlet boundary. In the cG case, the basis functions on the interface are identified accordingly to obtain a conforming subspace of V_D . For the remainder of this paper, we drop the superscript $X \in \{cG, dG\}$ and use the symbol V_h for both formulations. Depending on the considered formulation, one needs to use the right space V_h^X , $X \in \{cG, dG\}$. The IgA schemes (4) and (6) are equivalent to the system of linear IgA equations

$$\mathbf{K}\mathbf{u} = \mathbf{f}, \quad (10)$$

where $\mathbf{K} = (\mathbf{K}_{i,j})_{i,j \in \mathcal{I}_0}$, $\mathbf{f} = (\mathbf{f}_i)_{i \in \mathcal{I}_0}$ denote the stiffness matrix and the load vector, respectively, with $\mathbf{K}_{i,j} = a(N_{j,p}, N_{i,p})$ or $\mathbf{K}_{i,j} = a_h(N_{j,p}, N_{i,p})$ and $\mathbf{f}_i = \langle F, N_{i,p} \rangle$, and \mathbf{u} is the vector representation of u_h .

4 IETI-DP methods and their implementation

In this section, we recall the main ingredients for the cG-IETI-DP and dG-IETI-DP method. We focus mainly on the implementation, since this is the relevant part for parallelization.

4.1 Derivation of the method

A rigorous and formal definition of the cG-IETI-DP and dG-IETI-DP method is quite technical and not necessary for the parallelization, which is the purpose of this paper. Therefore, we are not going to present the whole derivation of each method. We will give a general description, which is valid for both methods. For a detailed derivation, we refer to [19] and [20].

The first step is to introduce additional dofs on the interface to decouple the local problems and incorporate their connection via Lagrange multipliers $\boldsymbol{\lambda}$. This is quite straightforward in the case of the cG formulation, but more involved in the dG case. In any of the two cases, we can equivalently rewrite (10) as: Find $(u, \boldsymbol{\lambda}) \in V_{h,e} \times U$ such that

$$\begin{bmatrix} K_e & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}, \quad (11)$$

where $V_{h,e} \supset V_h$, is the decoupled space with additional dofs and U is the set of Lagrange multipliers. The jump operator B enforces the “continuity” of the solution u in the sense that $\ker(B) \equiv V_h$. The matrix K_e is the block diagonal matrix of the patch local stiffness

matrices $K^{(k)}$, i.e., $K_e = \text{diag}(K^{(k)})$. Since B only acts on the patch interface dofs, we first can reorder the stiffness matrix in the following way

$$K^{(k)} = \begin{bmatrix} K_{BB}^{(k)} & K_{BI}^{(k)} \\ K_{IB}^{(k)} & K_{II}^{(k)} \end{bmatrix}, \quad f^{(k)} = \begin{bmatrix} f_B^{(k)} \\ f_I^{(k)} \end{bmatrix}$$

and then consider only the Schur complement representation: Find $(u_B, \boldsymbol{\lambda}) \in W \times U$ such that

$$\begin{bmatrix} S_e & B_B^T \\ B_B & 0 \end{bmatrix} \begin{bmatrix} u_B \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} g \\ 0 \end{bmatrix}, \quad (12)$$

where $S_e = \text{diag}(S_e^{(k)})$, $S_e^{(k)} = K_{BB}^{(k)} - K_{BI}^{(k)}(K_{II}^{(k)})^{-1}K_{IB}^{(k)}$ and $g^{(k)} = f_B - K_{BI}^{(k)}(K_{II}^{(k)})^{-1}f_I^{(k)}$. The space W is the restriction of $V_{h,e}$ to the interface. For completeness, we denote its “continuous” representation as \widehat{W} , i.e., $\ker(B_B) = \widehat{W}$. Equation (12) is valid for the cG-IETI-DP and dG-IETI-DP method, but the matrix K_e has difference entries and the number of boundary dofs (subscript B) is different. Fortunately, this does not change the way how the algorithm is implemented and parallelized. In the following, we will drop the subscript B in u_B and B_B for notational simplicity.

The matrix S_e is not invertible and, hence, we cannot build the Schur complement system of (12). To overcome this, we introduce an intermediate space \widetilde{W} , such that $\widehat{W} \subset \widetilde{W} \subset W$, and S_e restricted to \widetilde{W} , denoted by \widetilde{S} , is invertible. We introduce primal variables as a set $\Psi \subset \widehat{W}^*$ and define the spaces

$$\widetilde{W} := \{w \in W : \psi(w^{(k)}) = \psi(w^{(l)}), \forall \psi \in \Psi, \forall k > l\}$$

and

$$W_\Delta := \prod_{k=1}^N W_\Delta^{(k)}, \quad \text{with} \quad W_\Delta^{(k)} := \{w^{(k)} \in W^{(k)} : \psi(w^{(k)}) = 0 \forall \psi \in \Psi\}.$$

Moreover, we introduce the space $W_\Pi \subset \widehat{W}$ such that $\widetilde{W} = W_\Pi \oplus W_\Delta$. We call W_Π *primal space* and W_Δ *dual space*. Typically, the set Ψ corresponds to “continuous” vertex values, edge averages and/or face averages.

Since $\widetilde{W} \subset W$, there is a natural embedding $\widetilde{I} : \widetilde{W} \rightarrow W$. Let the jump operator restricted to \widetilde{W} be $\widetilde{B} := B\widetilde{I} : \widetilde{W} \rightarrow U^*$. Now we are in the position to reformulate problem (12) in the space \widetilde{W} as follows: Find $(u, \boldsymbol{\lambda}) \in \widetilde{W} \times U$:

$$\begin{bmatrix} \widetilde{S} & \widetilde{B}^T \\ \widetilde{B} & 0 \end{bmatrix} \begin{bmatrix} u \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} \widetilde{g} \\ 0 \end{bmatrix}, \quad (13)$$

where $\widetilde{g} := \widetilde{I}^T g$, and $\widetilde{B}^T = \widetilde{I}^T B^T$. Here, $\widetilde{I}^T : W^* \rightarrow \widetilde{W}^*$ denotes the adjoint of \widetilde{I} .

By construction, \widetilde{S} is SPD on \widetilde{W} . Hence, we can define the Schur complement F and the corresponding right-hand side as follows:

$$F := \widetilde{B}\widetilde{S}^{-1}\widetilde{B}^T, \quad d := \widetilde{B}\widetilde{S}^{-1}\widetilde{g}.$$

Hence, the saddle point system (13) is equivalent to the Schur complement problem:

$$\text{Find } \boldsymbol{\lambda} \in U : \quad F\boldsymbol{\lambda} = d. \quad (14)$$

Equation (14) is solved by means of the PCG algorithm, but it requires an appropriate preconditioner in order to obtain an efficient solver.

Recalling the definition of $S_e = \text{diag}(S_e^{(k)})_{k=1}^N$, we define the scaled Dirichlet preconditioner $M_{sD}^{-1} := B_D S_e B_D^T$, where B_D is a scaled version of the jump operator B . The scaled jump operator B_D is defined such that the operator enforces the constraints

$$\delta_j^{\dagger(l)}(\mathbf{u}^{(k)})_i^{(k)} - \delta_i^{\dagger(k)}(\mathbf{u}^{(l)})_j^{(k)} = 0 \quad \forall (i, j) \in B_e(k, l), \quad \forall l \in \mathcal{I}_{\mathcal{F}}^{(k)},$$

and

$$\delta_j^{\dagger(l)}(\mathbf{u}^{(k)})_i^{(l)} - \delta_i^{\dagger(k)}(\mathbf{u}^{(l)})_j^{(l)} = 0 \quad \forall (i, j) \in B_e(l, k), \quad \forall l \in \mathcal{I}_{\mathcal{F}}^{(k)},$$

where, for $(i, j) \in B_e(k, l)$, $\delta_i^{\dagger(k)} := \rho_i^{(k)} / \sum_{l \in \mathcal{I}_{\mathcal{F}}^{(k)}} \rho_j^{(l)}$ is an appropriate scaling. One can show, that the preconditioned system has a quasi-optimal condition number bound with respect to $H/h := \max_k (H_k/h_k)$, i.e.,

$$\kappa(M_{sD}^{-1} F|_{\ker(\tilde{B}^T)}) \leq C(1 + \log(H/h))^2, \quad (15)$$

for both versions, see [19], [18] and [5]. Moreover, numerical examples show also robustness with respect to jumps in the diffusion coefficient and only a weak dependence on the B-Spline degree p , see, e.g., [20], [19] and [5].

4.2 Implementation of the algorithm

Since F is symmetric and positive definite on \tilde{U} , we can solve the linear system $F\boldsymbol{\lambda} = d$ by means of the PCG algorithm, where we use M_{sD}^{-1} as preconditioner. The PCG does not require an explicit representation of the matrices F and M_{sD}^{-1} , since we just need their application to a vector. There are different ways to provide an efficient implementation. We will follow the concept of the energy minimizing primal subspaces. The idea is to split the space \tilde{W} into $\tilde{W}_{\Pi} \oplus \prod \tilde{W}_{\Delta}^{(k)}$, such that $\tilde{W}_{\Delta}^{(k)} \perp_S \tilde{W}_{\Pi}$ for all k , i.e., we choose $\tilde{W}_{\Pi} := \tilde{W}_{\Delta}^{\perp_S}$, see, e.g., [43] and [9]. By means of this choice, the operators \tilde{S} and \tilde{S}^{-1} have the following forms

$$\tilde{S} = \begin{bmatrix} S_{\Pi\Pi} & 0 \\ 0 & S_{\Delta\Delta} \end{bmatrix} \quad \text{and} \quad \tilde{S}^{-1} = \begin{bmatrix} S_{\Pi\Pi}^{-1} & 0 \\ 0 & S_{\Delta\Delta}^{-1} \end{bmatrix},$$

where $S_{\Pi\Pi}$ and $S_{\Delta\Delta}$ are the restrictions of \tilde{S} to the corresponding subspaces. We note that $S_{\Delta\Delta}$ can be seen as a block diagonal operator, i.e., $S_{\Delta\Delta} = \text{diag}(S_{\Delta\Delta}^{(k)})$.

The application of F and M_{sD}^{-1} is summarized in Algorithm 1.

4.2.1 Constructing a basis for the primal subspace First we need to provide an appropriate local basis $\{\tilde{\phi}_j\}_j^{n_{\Pi}}$ for \tilde{W}_{Π} , where n_{Π} is the number of primal variables. We request from the basis that it has to be nodal with respect to the primal variables, i.e., $\psi_i(\tilde{\phi}_j) = \delta_{i,j}$, for $i, j \in \{1, \dots, n_{\Pi}\}$. In order to construct such a basis, we introduce the constraint matrix $C^{(k)} : W^{(k)} \rightarrow \mathbb{R}^{n_{\Pi}^{(k)}}$ for each patch $\Omega^{(k)}$ which realizes the primal variables, i. e., $(C^{(k)}v)_j = \psi_{i(k,j)}(v)$ for $v \in W$ and $j \in \{1, \dots, n_{\Pi}^{(k)}\}$, where $n_{\Pi}^{(k)}$ is the number of primal variables associated with $\Omega^{(k)}$ and $i(k, j)$ the global index of the j -th

primal variable on $\Omega^{(k)}$. For each patch k , the basis functions $\{\tilde{\phi}_j^{(k)}\}_{j=1}^{n_\Pi^{(k)}}$ of $\widetilde{W}_\Pi^{(k)}$ are the solution of the system

$$\begin{bmatrix} K_{BB}^{(k)} & K_{BI}^{(k)} & C^{(k)T} \\ K_{IB}^{(k)} & K_{II}^{(k)} & 0 \\ C^{(k)} & 0 & 0 \end{bmatrix} \begin{bmatrix} \tilde{\phi}_j^{(k)} \\ \cdot \\ \tilde{\mu}_j^{(k)} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \mathbf{e}_j^{(k)} \end{bmatrix}, \quad (16)$$

where $\mathbf{e}_j^{(k)} \in \mathbb{R}^{n_\Pi^{(k)}}$ is the j -th unit vector. Here we use an equivalent formulation with the system matrix $K^{(k)}$. For each patch k , the LU factorization of this matrix is computed and stored.

Application of $S_{\Delta\Delta}^{(k)-1}$: The application of $S_{\Delta\Delta}^{(k)-1}$ corresponds to solving a local Neumann problem in the space \widetilde{W}_Δ , i.e., $S^{(k)}w^{(k)} = f_\Delta^{(k)}$ with the constraint $C^{(k)}w^{(k)} = 0$. This problem can be rewritten as a saddle point problem in the form

$$\begin{bmatrix} K_{BB}^{(k)} & K_{BI}^{(k)} & C^{(k)T} \\ K_{IB}^{(k)} & K_{II}^{(k)} & 0 \\ C^{(k)} & 0 & 0 \end{bmatrix} \begin{bmatrix} w^{(k)} \\ \cdot \\ \cdot \end{bmatrix} = \begin{bmatrix} f_\Delta^{(k)} \\ 0 \\ 0 \end{bmatrix}.$$

From (16), the LU factorization of the matrix is already available.

Application of $\mathbf{S}_{\Pi\Pi}^{(k)-1}$: The matrix $\mathbf{S}_{\Pi\Pi}$ can be assembled from the patch local matrices $\mathbf{S}_{\Pi\Pi}^{(k)}$. Let $\{\tilde{\phi}_j^{(k)}\}_{j=1}^{n_\Pi^{(k)}}$ be the basis of $\widetilde{W}_\Pi^{(k)}$. The construction of $\{\tilde{\phi}_j^{(k)}\}_{j=1}^{n_\Pi^{(k)}}$ in (16) provides

$$\begin{aligned} \left(\mathbf{S}_{\Pi\Pi}^{(k)}\right)_{i,j} &= \left\langle S^{(k)}\tilde{\phi}_i^{(k)}, \tilde{\phi}_j^{(k)} \right\rangle = - \left\langle C^{(k)T}\tilde{\mu}_i^{(k)}, \tilde{\phi}_j^{(k)} \right\rangle = - \left\langle \tilde{\mu}_i^{(k)}, C^{(k)}\tilde{\phi}_j^{(k)} \right\rangle \\ &= - \left\langle \tilde{\mu}_i^{(k)}, \mathbf{e}_j \right\rangle^{(k)} = - \left(\tilde{\mu}_i^{(k)}\right)_j, \end{aligned}$$

where $i, j \in \{1, \dots, n_\Pi^{(k)}\}$. Therefore, we can reuse the Lagrange multipliers $\tilde{\mu}_i^{(k)}$ obtained in (16), and can assemble $\mathbf{S}_{\Pi\Pi}^{(k)}$ from them. Once $\mathbf{S}_{\Pi\Pi}$ is assembled, the LU factorization can be calculated and stored.

4.2.2 Application of \tilde{I} and \tilde{I}^T The last building block is the embedding $\tilde{I} : \widetilde{W} \rightarrow W$ and its adjoint $\tilde{I}^T : W^* \rightarrow \widetilde{W}^*$. Recall the direct splitting $W^{(k)} = W_\Delta^{(k)} \oplus W_\Pi^{(k)}$. Let us denote by $\Phi^{(k)} = [\tilde{\phi}_1^{(k)}, \dots, \tilde{\phi}_{n_\Pi^{(k)}}^{(k)}]$ the coefficient representation of the basis for $W_\Pi^{(k)}$.

Given the primal part w_Π of a function in \widetilde{W} , we obtain its restriction to $\widetilde{W}_\Pi^{(k)}$ via an appropriately defined restriction matrix $\mathbf{R}^{(k)}$, i.e. $w_\Pi^{(k)} = \mathbf{R}^{(k)}w_\Pi$. The corresponding function is then given by $w_\Pi^{(k)} = \Phi^{(k)}\mathbf{R}^{(k)}w_\Pi$.

Following the lines in [43], we can formulate the operator $\tilde{I} : \widetilde{W} \rightarrow W$ as

$$\begin{bmatrix} w_\Pi \\ w_\Delta \end{bmatrix} \mapsto w := \Phi \mathbf{R} w_\Pi + w_\Delta,$$

where Φ and \mathbf{R} are block versions of $\Phi^{(k)}$ and $\mathbf{R}^{(k)}$, respectively. The second function is its adjoint operation $\tilde{I}^T : W^* \rightarrow \widetilde{W}^*$. It can be realized in the following way

$$f \mapsto \begin{bmatrix} \mathbf{f}_\Pi \\ f_\Delta \end{bmatrix} = \begin{bmatrix} \mathbf{A}\Phi^T \mathbf{f} \\ f - C^T \Phi^T f \end{bmatrix},$$

where \mathbf{A} is the corresponding assembling operator to \mathbf{R} , i.e., $\mathbf{A} = \mathbf{R}^T$. A more extensive discussion and derivation can be found in [43].

Algorithm 1 Algorithm for calculating $\boldsymbol{\nu} = F\boldsymbol{\lambda}$ and $\boldsymbol{\nu} = M_{sD}^{-1}\boldsymbol{\lambda}$ for given $\boldsymbol{\lambda} \in U$

```

procedure  $F(\boldsymbol{\lambda})$ 
  Application of  $B^T : \{f^{(k)}\}_{k=1}^N = B^T \boldsymbol{\lambda}$ 
  Application of  $\tilde{I}^T : \{\mathbf{f}_\Pi, \{f_\Delta^{(k)}\}_{k=1}^N\} = \tilde{I}^T(\{f^{(k)}\}_{k=1}^N)$ 
  Application of  $\tilde{S}^{-1}$  :
  Begin
     $\mathbf{w}_\Pi = \mathbf{S}_{\Pi\Pi}^{-1} \mathbf{f}_\Pi$ 
     $w_\Delta^{(k)} = S_{\Delta\Delta}^{(k)-1} f_\Delta^{(k)} \quad \forall k = 1, \dots, N$ 
  End
  Application of  $\tilde{I} : \{w^{(k)}\}_{k=1}^N = \tilde{I}(\{\mathbf{w}_\Pi, \{w_\Delta^{(k)}\}_{k=1}^N\})$ 
  Application of  $B : \boldsymbol{\nu} = B(\{w^{(k)}\}_{k=1}^N)$ 
end procedure
procedure  $M_{sD}^{-1}(\boldsymbol{\lambda})$ 
  Application of  $B_D^T : \{w^{(k)}\}_{k=1}^N = B_D^T \boldsymbol{\lambda}$ 
  Application of  $S_e$  :
  Begin
    Solve  $K_{II}^{(k)} x^{(k)} = -K_{IB}^{(k)} w^{(k)} \quad \forall k = 1, \dots, N$ 
     $v^{(k)} = K_{BB}^{(k)} w^{(k)} + K_{BI}^{(k)} x^{(k)} \quad \forall k = 1, \dots, N$ 
  End
  Application of  $B_D : \boldsymbol{\nu} = B_D(\{v^{(k)}\}_{k=1}^N)$ 
end procedure

```

5 Parallelization of the building blocks

Here we investigate how the single operations can be executed in parallel in a distributed memory setting. The parallelization of the method is performed with respect to the patches, i.e., one or several patches are assigned to a processor. The required communication has to be understood as communication between patches, which are assigned to different processors. The majority of the used MPI methods are performed in its non-blocking version. We aim at overlapping computations with communications wherever possible.

5.1 Parallel version of PCG

We solve $F\boldsymbol{\lambda} = d$ with the preconditioned CG method. This requires a parallel implementation of CG, where we follow the approach presented in Section 2.2.5.5 in [43], see also [11]. This approach is based on the concept of accumulated and distributed vectors. We say a vector $\boldsymbol{\lambda}_{acc} = [\boldsymbol{\lambda}_{acc}^{(q)}]$ is an *accumulated* representation of $\boldsymbol{\lambda}$, if $\boldsymbol{\lambda}_{acc}^{(q)}(k_q(i)) = \boldsymbol{\lambda}(i)$, where i is the global index corresponding to the local index $k_q(i)$ on processor q . On the

contrary, $\lambda_{dist} = [\lambda_{dist}^{(q)}]$ is a *distributed* representation of λ , if the sum of all processor local contributions give the global vector, i.e., $\lambda_{dist}(i) = \sum_q \lambda_{dist}^{(q)}(k_q(i))$. Hence, each processor only holds the part of λ , which belongs to its patches, either in a distributed or accumulated description. The Lagrange multipliers and the search direction of the CG are represented in the accumulated setting, whereas the residual is given in the distributed representation. In order to achieve the accumulated representation, information exchange between the neighbours of a patch is required. This is done after applying the matrix and the preconditioner, respectively and implemented via `MPI_Send` and `MPI_Recv` operations.

The last aspect in the parallel CG implementation is the realization of scalar products. Given a distributed representation u_{dist} of u and an accumulated representation of v_{acc} of v , the scalar product $(u, v)_{l^2}$ is then given by $(u, v)_{l^2} = \sum_q (u_{dist}^{(q)}, v_{acc}^{(q)})_{l^2}$, i.e., first the local scalar products are formed, globally added, and distributed with `MPI_Allreduce`.

5.2 Assembling

The assembling routine of the IETI-DP algorithm consists of the following steps:

1. Assemble the patch local stiffness matrices and right hand side,
2. assemble the system matrix in (16) and calculating its LU-factorization,
3. assemble $S_{\Pi\Pi}$ and calculating its LU-factorization,
4. calculate the LU-factorization of $K_{II}^{(k)}$,
5. calculate the right hand side $\{g_{\Pi}, g_{\Delta}\} = \tilde{g} \in \widetilde{W}^*$, with $g^{(k)} = f_B - K_{BI}^{(k)}(K_{II}^{(k)})^{-1}f_I^{(k)}$.

Most of the tasks are completely independent of each other and, hence, can be performed in parallel. Only the calculation of $S_{\Pi\Pi}$ and $\tilde{g} = \tilde{I}^T g$ require communication, which will be handled in Section 5.3.

The LU-factorization of $S_{\Pi\Pi}$ is only required at one processor, since it has to be solved only once per CG iteration. According to [30], it is advantageous to distribute this matrix to all other processors in order to reduce communication in the solver part, see [33] and references therein for improving scalability based on a different approach. In the current paper, we investigate cases, where one, several and all processor hold the LU-factorization of $S_{\Pi\Pi}$. Therefore, each processor is assigned to exactly one holder of $S_{\Pi\Pi}$. This relation is implemented by means of an additional MPI communicator.

We note that, for extremely large scale problems with $\geq 10^5$ subdomains, one has to consider different strategies dealing with $S_{\Pi\Pi}^{-1}$. Most commonly one uses AMG and solves $S_{\Pi\Pi}u_{\Pi} = f_{\Pi}$ in an inexact way, see, e.g., [28] and [32]. When considering a moderate number of patches, i.e., $10^3 - 10^4$, the approach using the LU-factorization of $S_{\Pi\Pi}$ is the most efficient one. In this paper, we restrict ourselves to this case.

The patch local matrix $S_{\Pi\Pi}^{(k)}$ is obtained as a part of the solutions of (16) and the assembling of the global matrix $S_{\Pi\Pi}$ is basically a `MPI_gatherv` operation. In the case where all processors hold $S_{\Pi\Pi}$ we use `MPI_allgatherv`. If several processors hold the LU factorization, we just call `MPI_gatherv` on each of these processors. A different possibility would be to first assemble $S_{\Pi\Pi}$ on one patch, distribute it to the other holders and then calculate the LU-factorization on each of the processors.

5.3 Solver and Preconditioner

More communication is involved in the solver part. According to Algorithm 1, we have to perform the following operations:

1. application of B and B^T and its scaled versions
2. application of \tilde{I} and \tilde{I}^T
3. application of \tilde{S}^{-1}
4. application of S^{-1}

The only operations which require communication are \tilde{I} and \tilde{I}^T . To be more precise, the communication is hidden in the operators \mathbf{A} and \mathbf{R} , see Section 4.2, all other operations are block operations, where the corresponding matrices are stored locally on each processor. In principle, their implementation is given by accumulating and distributing values. The actual implementation depends on how many processors hold the coarse grid problem.

In order to implement \tilde{I} , we need the distribution operation \mathbf{R} . If all processors hold $S_{\Pi\Pi}$, this operations reduces to just extracting the right entries. Hence it is local and no communication is required. Otherwise, each holder of $S_{\Pi\Pi}$ reorders and duplicates the entries of \mathbf{w}_{Π} in such a way, that all entries corresponding to the patches of a single slave are in a contiguous block of memory. Then we utilize the `MPI_scatter` method to distribute only the necessary data to all slave processors. See Figure 1 for an illustration.

We arrive at the implementation of \tilde{I}^T . Each processor stores the values of $\mathbf{w}_{\Pi}^{(k)}$ in a vector $\tilde{\mathbf{w}}_{\Pi}^{(k)}$ of length n_{Π} already in such a way, that $\sum_{k=1}^N \tilde{\mathbf{w}}_{\Pi}^{(k)} = \mathbf{w}_{\Pi}$. Storing the entries in this way enables the use the MPI reduction operations to efficiently assemble the local contributions. If only one processor holds the coarse problem, we use the `MPI_Reduce` method to perform this operation. Similarly, if all processors hold $S_{\Pi\Pi}$, we utilize the `MPI_Allreduce` method. If several processors have the coarse grid problem, we use a two level approach. First, each master processor collects the local contributions from its slaves using the `MPI_Reduce` operation. In the second step, all the master processors perform an `MPI_Allreduce` operation to accumulate the contributions from each group and simultaneously distribute the result under them. This procedure is visualized in Figure 1.

6 Numerical examples

We consider the model problem (1) in the two dimensional computational domain $\Omega = (0, 1)^2$ formed by $32 \times 32 = 1024$ patches. Each of them is a square arranged in a uniform grid. For the three dimensional case we consider the domain $\Omega = (0, 1)^2 \times (0, 2)$, partitioned into $8 \times 8 \times 16$ regular cubes. Note that, in IgA framework, we cannot choose the number of subdomains as freely as in the finite element case since they are fixed by the geometry. Therefore, the number of 1024 subdomains stays constant throughout the tests. Since we are interested in the parallel scalability of the proposed algorithms, we assume for simplicity homogeneous diffusion coefficients $\alpha \equiv 1$. In all tests we consider the smooth right hand side $f(x, y) = 20\pi^2 \sin(4\pi(x + 0.4)) \sin(2\pi(y + 0.3))$, corresponding to the exact solution $u(x, y) = \sin(4\pi(x + 0.4)) \sin(2\pi(y + 0.3)) + x + y$. For the discretization, we use

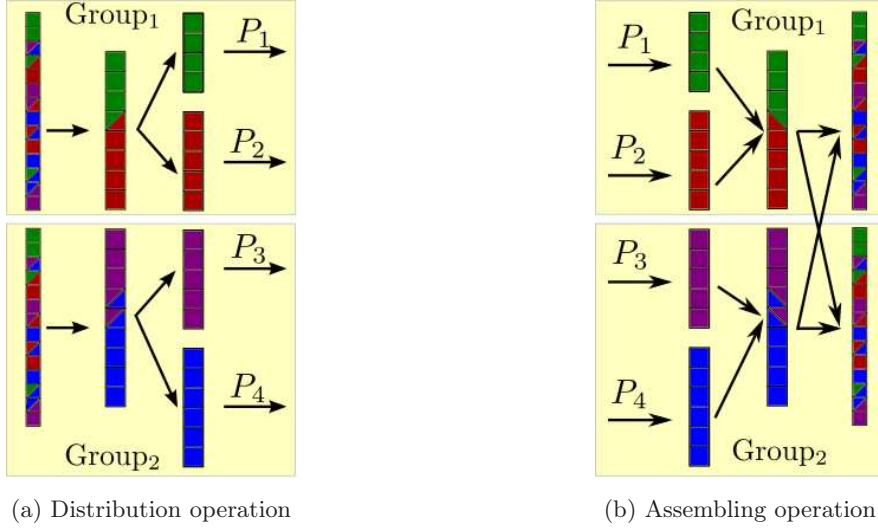


Fig. 1: Distribution and assembling operation, illustrated for four processors, partitioned into two groups corresponding to two S_{III}^{-1} holder.

tensor B-Spline spaces V_h of different degree p . We increase the B-Spline degree in such a way that the number of knots stay the same, i.e., the smoothness of V_h increases.

We investigate the scaling behaviour of the cG-IETI-DP and dG-IETI-DP method. Although, we consider also the dG variant, we restrict ourselves to matching meshes. Otherwise, it would not be possible to compare the two methods. Moreover, some patches would have a significant larger number of dofs, which leads to load imbalances and affects the scaling in a negative way. The domain is refined in a uniform way by inserting a single knot for each dimension on each knot span. We denote by H_k the patch diameter and by h_k the characteristic mesh size on $\Omega^{(k)}$. The set of primal variables is chosen by continuous patch vertices and interface averages for the two dimensional setting. For the three dimensional examples, we choose only continuous edge averages in order to keep the number of primal variables small.

The preconditioned conjugate gradient method is used to solve (14) with the scaled Dirichlet preconditioner M_{sD}^{-1} . We choose zero initial guess and a relative reduction of the residual of 10^{-8} . For solving the local systems and the coarse grid problem, a direct solver is used.

The algorithm is realized in the isogeometric open source C++ library G+SMO [41], which is based on the Eigen library [17]. We utilize the PARDISO 5.0.0 Solver Project [36] for performing the LU factorizations. The code is compiled with the gcc 4.8.3 compiler with optimization flag -O3. For the communication between the processors, we use the MPI 2 standard with the OpenMPI 1.10.2 implementation. The results are obtain on the RADON1 cluster at Linz. We use 64 out of 68 available nodes, each equipped with 2x Xeon E5-2630v3 “Haswell” CPU (8 Cores, 2.4Ghz, 20MB Cache) and 128 GB RAM. This gives the total number of 1024 available cores.

We investigate two quantities, the assembling phase and the solving phase. In the assembling phase, we account for the time used for

- assembling the local matrices and right hand sides,
- LU-factorization of K_{II} ,

- LU-factorization of $\begin{bmatrix} K & C^T \\ C & 0 \end{bmatrix}$,
- calculation of $\tilde{\Phi}$ and $\tilde{\mu}$,
- assembling the coarse grid matrix $S_{\Pi\Pi}$ and calculation of its LU factorization.

As already indicated in Section 5, $S_{\Pi\Pi}$ is only assembled on certain processors. The solving phase consists of the CG algorithm for (14) and the back-substitution to obtain the solution from the Lagrange multipliers. The main ingredients are

- application of F ,
- application of M_{sD}^{-1} .

In Section 6.1 and Section 6.2, we study the weak and strong scaling behaviour for the cG-IETI-DP and the dG-IETI-DP method. In this two sections, we assume that only one processor holds the coarse grid matrix $S_{\Pi\Pi}$. The comparison of having a different number of $S_{\Pi\Pi}$ holders is done in Section 6.3.

6.1 Weak scaling

In this subsection we investigate the weak scaling behaviour, i.e., the relation of problem size and number of processors is constant. In each refinement step we multiply the number of used cores by $2^d, d \in \{2, 3\}$. The ideal behaviour would be a constant time for each refinement.

First, we consider the two dimensional case. We apply three initial refinements and start with a single processor and perform up to additional 5 refinements with maximum 1024 processors. We choose as primal variables continuous vertex values and edge averages. The results for degree $p \in \{2, 3, 4\}$ are illustrated in Figure 2. The first row of figures corresponds to the cG-IETI-DP method, and the second one corresponds to the dG-IETI-DP method. The left column of Table 1 summarizes timings and the speedup for the cG-IETI-DP method, whereas the right column presents the results for the dG-IETI-DP method. For each method, we investigate the weak scaling for the assembling and solution phase. As in Figure 2, we present the scaling and timings for $p \in \{2, 3, 4\}$.

We observe that the time used for the assembling stays almost constant, hence shows quite optimal behaviour. However, the time for solving the system increases when refining and increasing the number of used processors. Especially, when considering the largest number of processors, we see a clear increase of the solution time. One reason is that the number of iterations slightly increases when increasing the system size. This is due to the quasi optimal condition number bound of the IETI-DP type methods, cf. (15). Secondly, as already pointed out in Section 5, the solving phase consists of more communication between processors, which cannot be completely overlapped with computations. Moreover, one also has to take in account global synchronization points in the conjugate gradient method.

Next, we consider the weak scaling for the three dimensional case. As already indicated in the introduction of this section, we choose only continuous edge averages as primal variables. We perform the tests in the same way as for the two dimensional case. However, we already start with two processors and perform two initial refinements. Multiplying

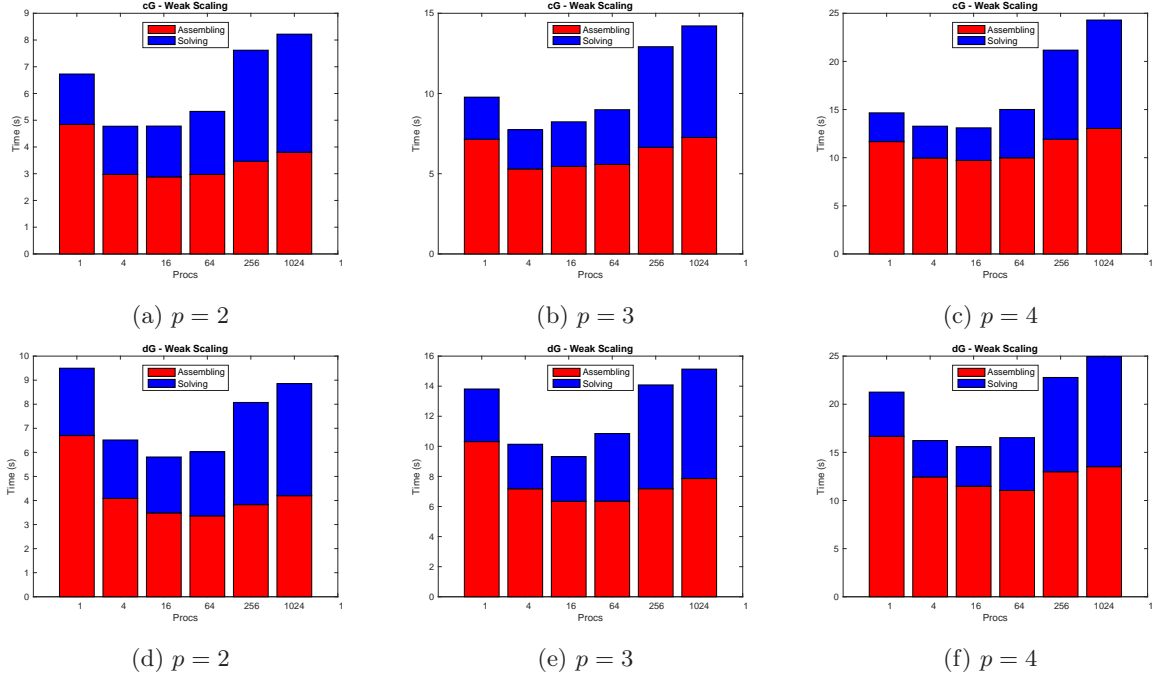


Fig. 2: Weak scaling of the cG-IETI-DP (first row) and dG-IETI-DP (second row) method for B-Spline degrees $p \in \{2, 3, 4\}$ in two dimensions. Each degree corresponds to one column.

the number of used processors by 8 with each refinement, we end up again with 1024 processors on the finest grid. The two algorithms behave similar to the two dimensional case, where the assembling phase gives quite good results and the solver phase shows again an increasing time after each refinement. The results are visualized in Figure 3 and summarized in Table 2. Note, for the dG-IETI-DP method with $p = 4$ and ~ 54 Mio. dofs, we exceeded the memory capacity of the cluster.

6.2 Strong scaling

Secondly, we are investigating the strong scaling behaviour. Now we fix the problem size and increase the number of processors. In the optimal case, the time used by a certain quantity reduces in the same way as the number of used processors increases. We use the same primal variables for the strong scaling studies as in the weak scaling studies in Section 6.1.

Again as in Section 6.1, we begin with the two dimensional example. We perform 7 initial refinements and end up with 17 Mio. dofs on 1024 subdomains. We start already with 4 processors in the initial case and do 8 refinements until we reach 1024 cores. Similar to Section 6.1, the results for $p \in \{2, 3, 4\}$ are illustrated in Figure 4 and summarized in Table 3.

We observe that the assembling phase has a quite good scaling performance, as already observed for the weak scaling results in Section 6.1. Moreover, the higher the B-Spline degree, the better the parallel performance behaves. This holds due to increased computational costs for the parallel part. Similar to the weak scaling results, the solver phase does not provide such an excellent scaling as the assembling phase. Still, we obtain a scaling

	cG-IETI-DP		$p = 2$			dG-IETI-DP		$p = 2$		
# procs	#dofs	Iter.	Ass. Time	Solv. Time	Total Time	#dofs	Iter.	Ass. Time	Solv. Time	Total Time
1	99104	7	4.8	1.9	6.7	133824	8	6.7	2.8	9.5
4	328224	8	3.0	1.8	4.8	394688	9	4.1	2.4	6.5
16	1179680	9	2.9	1.9	4.8	1309632	10	3.5	2.3	5.8
64	4455456	10	3.0	2.4	5.4	4712384	11	3.4	2.7	6.1
256	17298464	11	3.5	4.2	7.7	17809344	11	3.8	4.2	8.0
1024	68150304	11	3.8	4.4	8.2	69169088	12	4.2	4.7	8.9

	cG-IETI-DP		$p = 3$			dG-IETI-DP		$p = 3$		
# procs	#dofs	Iter.	Ass. Time	Solv. Time	Total Time	#dofs	Iter.	Ass. Time	Solv. Time	Total Time
1	120576	8	7.2	2.6	9.8	159264	8	10.3	3.5	13.8
4	366080	9	5.3	2.4	7.7	436512	9	7.2	3.0	10.2
16	1250304	10	5.5	2.8	8.3	1384224	10	6.3	3.0	9.3
64	4591616	10	5.6	3.4	9.0	4852512	11	6.4	4.5	10.9
256	17565696	11	6.6	6.3	12.9	18080544	12	7.2	6.9	14.1
1024	68679680	12	7.3	7.0	14.3	69702432	12	7.9	7.3	15.2

	cG-IETI-DP		$p = 4$			dG-IETI-DP		$p = 4$		
# procs	#dofs	Iter.	Ass. Time	Solv. Time	Total Time	#dofs	Iter.	Ass. Time	Solv. Time	Total Time
1	144096	8	11.7	3.0	14.7	186752	9	16.7	4.6	21.3
4	405984	9	10.0	3.3	13.3	480384	10	12.4	3.8	16.2
16	1322976	10	9.7	3.4	13.1	1460864	11	11.5	4.1	15.6
64	4729824	11	10.0	5.0	15.0	4994688	11	11.0	5.5	16.5
256	17834976	12	11.9	9.3	21.2	18353792	12	13.0	9.8	22.8
1024	69211104	13	13.0	11.3	24.3	70237824	13	13.5	11.4	24.9

Table 1: Weak scaling results for the two dimensional testcase for the cG and dG IETI-DP method. Left column contains results for the cG variant and the right column for the dG version. Each row corresponds to a fixed B-Spline degree $p \in \{2, 3, 4\}$

from around 500 when using 1024 processors. We note that the degree of the B-Splines does not seem to have such a significant effect on the scaling for the solver phase as for the assembling phase.

In the three dimensional example we perform four initial refinements and obtain around 5 Mio. dofs. The presentation of the results is done in the same way as in the previous examples, see Figure 5 and Table 4. Also in three dimensions the cG-IETI-DP algorithms behaves very similarly to the two dimensional case, showing excellent scaling results. However, the dG version of the algorithm shows a good scaling but not as promising as cG version. Especially, when considering $p = 2$, we observe degraded scalability for the assembling phase. Having a closer look at the timings, we observe that this originates from small load imbalances in the interior domains, due to the additional layer of dofs and the larger number of primal variables. The latter one leads to an increased time in solving (16), due to a larger number of right hand sides on the interior subdomains. One can further optimize the three dimensional case, by considering different strategies for the primal variables, where one aims for smaller and more equally distributed numbers of primal variables.

6.3 Study on the number of $S_{\Pi\Pi}^{-1}$ holders

In this last section of the numerical experiments, we want to investigate the influence of the number of holders of $S_{\Pi\Pi}^{-1}$ on the scaling behaviour. As already indicated in Section 5.3,

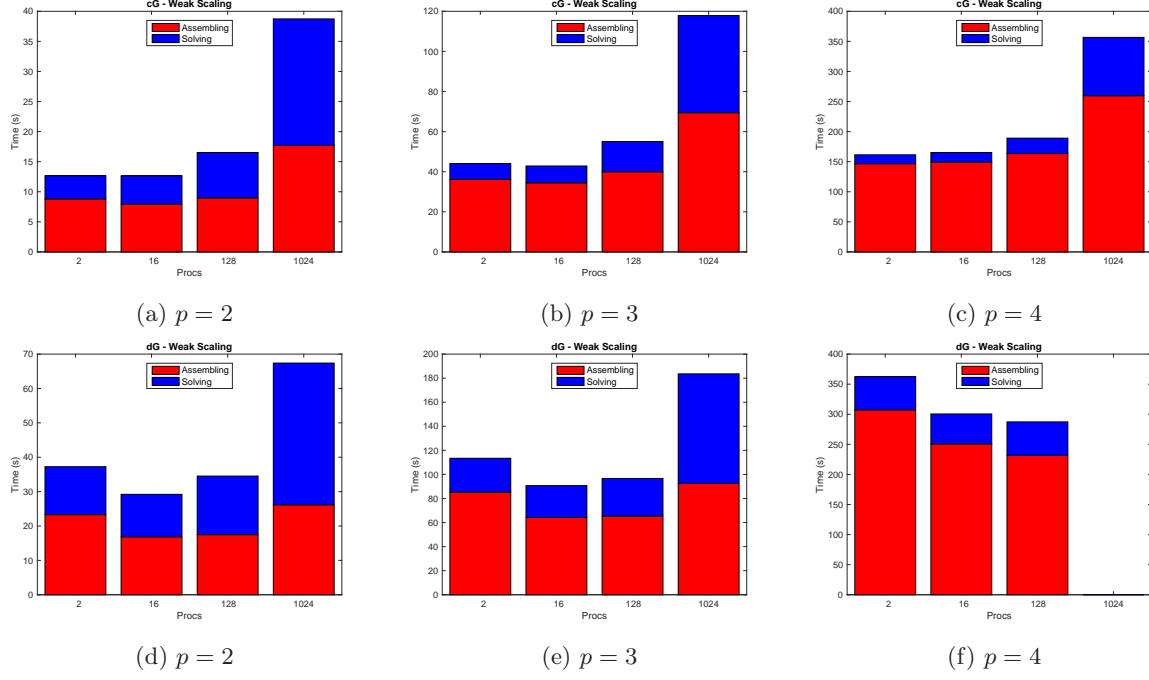


Fig. 3: Weak scaling of the cG-IETI-DP (first row) and dG-IETI-DP (second row) method for B-Spline degrees $p \in \{2, 3, 4\}$ in three dimensions. Each degree corresponds to one column. No timings are obtained in the case of 1024 cores in (f) due to memory limitations

if more processors hold the LU-factorization of the coarse grid matrix, it is possible to decrease the communication effort after applying $S_{\Pi\Pi}^{-1}$, while having more communication before the application. The advantage of this strategy is to be able to have a better overlap of communication with computations. However one has to take into account, that this also increases the communication in the assembling phase, since the local contribution $S_{\Pi\Pi}^{(k)}$ has to be sent to all the master processors.

We only consider the two dimensional domain, where we perform 7 initials refinements, but on a decomposition with 4096 subdomains and end up with around 70 Mio. dofs. This gives a comparable setting as in Section 6.1 having the most refined domain. In order to better observe the influence of the number of $S_{\Pi\Pi}^{-1}$ holders, we increase the number of subdomains, leading to a larger coarse grid problem. We only investigate the case of using 1024 processors and the number of $S_{\Pi\Pi}$ holders given by $2^n, n \in \{0, 1, \dots, 10\}$. Hence, we obtain the number of master processors ranging from 1 to 1024, such that each master has the same number of slaves. The results are summarized in Figure 6 and Table 5.

We observe that choosing several holders of the coarse grid problem in the cG version does not really have a significant effect. However, in the dG version, due to an increased number of primal variables, the use of several holders actually increases the performance of the solver by around 10%. Nevertheless, what is gained in the solving part does not pay off with the additional effort in the assembling phase. Considering the total computation time in Table 5, the best options is still either using only a single coarse grid problem on one processor or making a redundant factorization on each processor.

	cG-IETI-DP			$p = 2$			dG-IETI-DP			$p = 2$		
# procs	#dofs	Iter.	Ass. Time	Solv. Time	Total Time	#dofs	Iter.	Ass. Time	Solv. Time	Total Time		
2	220896	16	8.8	3.8	12.6	396932	26	23.3	14.0	37.3		
16	1023200	17	8.0	4.7	12.7	1551400	27	16.9	12.3	29.2		
128	5969376	17	9.0	7.5	16.5	7730288	28	17.5	17.0	34.5		
1024	40238048	19	17.7	21.0	38.7	46577920	28	26.2	41.2	67.4		

	cG-IETI-DP			$p = 3$			dG-IETI-DP			$p = 3$		
# procs	#dofs	Iter.	Ass. Time	Solv. Time	Total Time	#dofs	Iter.	Ass. Time	Solv. Time	Total Time		
2	350840	17	36.2	7.9	44.1	598405	29	85.3	28.1	113.4		
16	1361976	18	34.5	8.4	42.9	2005737	30	64.2	26.5	90.7		
128	7020728	18	40.0	15.0	55.0	8985265	30	65.4	31.3	96.7		
1024	43894200	21	69.4	48.4	117.8	50613825	31	92.6	91.0	183.6		

	cG-IETI-DP			$p = 4$			dG-IETI-DP			$p = 4$		
# procs	#dofs	Iter.	Ass. Time	Solv. Time	Total Time	#dofs	Iter.	Ass. Time	Solv. Time	Total Time		
2	523776	18	146.6	14.8	161.4	853878	32	307.3	55.1	362.4		
16	1768320	18	149.2	15.9	165.1	2538650	32	250.7	49.8	300.5		
128	8188800	20	163.6	25.6	189.2	10367970	34	232.0	55.4	287.4		
1024	47765376	22	259.7	96.9	356.6	~54000000	x	x	x	x		

Table 2: Weak scaling results for the three dimensional testcase for the cG and dG IETI-DP method. Left column contains results for the cG variant and the right column for the dG version. Each row corresponds to a fixed B-Spline degree $p \in \{2, 3, 4\}$. No timings are available for the dG-IETI-DP method with $p = 4$ on 1024 cores due to memory limitations.

7 Conclusion

We have investigated the parallel scalability of the cG-IETI-DP and dG-IETI-DP method, respectively. Numerical tests showed a very good scalability in the strong and weak scaling for the assembling phase for both methods. We reached a speedup of approximately 900 when using 1024 cores. Although the speedup of the solver phase is not as good as the one for the assembler phase, we still reached a speedup of around 500 when using 1024 cores. One can even increase the parallel performance of the solver part by increasing the number of processors, which are holding the coarse grid problem. However, numerical examples have shown that this does not really pay off in the total time, due to an increased assembling time. To summarize, we saw that the proposed methods are well suited for large scale parallelization of assembling and solving IgA equations in two and three dimensions.

Acknowledgements

This work was supported by the Austrian Science Fund (FWF) under the grant W1214, project DK4. This support is gratefully acknowledged. Moreover, the author wants to thank Prof. Ulrich Langer for the valuable comments and support during the preparation of the paper. He also gratefully acknowledges the support of Katharina Rafetseder of the Institute of Numerical Mathematics, Johannes Kepler University Linz (JKU) and Ioannis Touloupoulos and Angelos Mantzaflaris, Radon Institute of Computational and Applied Mathematics Linz (RICAM).




Fig. 4: Strong scaling of the cG-IETI-DP (left column) and dG-IETI-DP (right column) method for B-Spline degrees $p \in \{2, 3, 4\}$ in two dimensions. The markers $\{\circ, *, \diamond\}$ as well as different shades of red (assembling phase) and blue (solver phase) correspond to the degrees $\{2, 3, 4\}$.

References

References

1. Y. Bazilevs, L. Beirão da Veiga, J. A. Cottrell, T. J. R. Hughes, and G. Sangalli. Isogeometric analysis: Approximation, stability and error estimates for h -refined meshes. *Math. Models Methods Appl. Sci.*, 16(7):1031–1090, 2006.
2. Y. Bazilevs, V. Calo, J. Cottrell, J. Evans, T. Hughes, S. Lipton, M. Scott, and T. Sederberg. Isogeometric analysis using T-splines. *Computer Methods in Applied Mechanics and Engineering*, 199(5–8):229 – 263, 2010. Computational Geometry and Analysis.
3. L. Beirão da Veiga, A. Buffa, G. Sangalli, and R. Vázquez. Mathematical analysis of variational isogeometric methods. *Acta Numerica*, 23:157–287, 2014.
4. L. Beirão da Veiga, D. Cho, L. F. Pavarino, and S. Scacchi. Overlapping Schwarz methods for isogeometric analysis. *SIAM J. Numer. Anal.*, 50(3):1394–1416, 2012.
5. L. Beirão da Veiga, D. Cho, L. F. Pavarino, and S. Scacchi. BDDC preconditioners for isogeometric analysis. *Math. Models Methods Appl. Sci.*, 23(6):1099–1142, 2013.
6. L. Beirão da Veiga, D. Cho, L. F. Pavarino, and S. Scacchi. Isogeometric Schwarz preconditioners for linear elasticity systems. *Comput. Methods Appl. Mech. Eng.*, 253:439–454, 2013.
7. J. A. Cottrell, T. J. R. Hughes, and Y. Bazilevs. *Isogeometric Analysis, Toward Integration of CAD and FEA*. John Wiley and Sons, 2009.
8. D. A. Di Pietro and A. Ern. *Mathematical aspects of discontinuous Galerkin methods*. Berlin: Springer, 2012.
9. C. R. Dohrmann. A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci. Comput.*, 25(1):246–258, 2003.
10. C. R. Dohrmann. An approximate BDDC preconditioner. *Numerical Linear Algebra with Applications*, 14(2):149–168, 2007.
11. C. C. Douglas, G. Haase, and U. Langer. *Tutorial on Elliptic PDE Solvers and Their Parallelization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003.
12. M. Dryja, J. Galvis, and M. Sarkis. BDDC methods for discontinuous Galerkin discretization of elliptic problems. *J. Complexity*, 23(4-6):715–739, 2007.
13. M. Dryja, J. Galvis, and M. Sarkis. A FETI-DP preconditioner for a composite finite element and discontinuous Galerkin method. *SIAM J. Numer. Anal.*, 51(1):400–422, 2013.
14. M. Dryja and M. Sarkis. 3-d feti-dp preconditioners for composite finite element-discontinuous galerkin methods. In *Domain Decomposition Methods in Science and Engineering XXI*, pages 127–140. Springer, 2014.
15. C. Giannelli, B. Jüttler, and H. Speleers. THB-splines: the truncated basis for hierarchical splines. *Comput. Aided Geom. Design*, 29, 2012.
16. C. Giannelli, B. Jüttler, and H. Speleers. Strongly stable bases for adaptively refined multilevel spline spaces. *Advances in Computational Mathematics*, 40:459–490, 2014.

2d	$p = 2$				$p = 3$				$p = 4$			
cG-IETI-DP	assembling phase		solving phase		assembling phase		solving phase		assembling phase		solving phase	
# procs	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.
4	190.8	4	138.2	4	364.5	4	202.6	4	653.7	4	279.1	4
8	94.8	8	88.8	6	181.2	8	141.8	6	325.9	8	202.7	6
16	47.1	16	45.9	12	89.9	16	71.9	11	162.3	16	102.6	11
32	23.1	32	22.9	24	44.5	32	35.7	23	80.4	32	51.3	22
64	11.6	65	11.8	46	22.4	65	18.5	44	40.2	64	26.3	42
128	5.9	127	7.3	75	11.3	128	11.1	73	20.4	128	14.8	75
256	3.0	247	4.1	133	5.7	251	6.1	131	10.4	250	8.7	128
512	1.6	471	2.1	257	2.9	487	3.2	250	5.3	493	4.7	235
1024	0.9	819	1.1	472	1.6	891	1.6	494	2.8	917	2.4	456

dG-IETI-DP	assembling phase		solving phase		assembling phase		solving phase		assembling phase		solving phase	
# procs	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.
4	216.6	4	144.0	4	402.2	4	225.5	4	711.8	4	294.2	4
8	106.9	8	92.9	6	199.5	8	156.8	5	352.6	8	210.6	5
16	52.5	16	47.9	12	98.1	16	80.4	11	174.3	16	106.4	11
32	25.1	34	25.7	22	47.5	33	42.2	21	84.9	33	55.2	21
64	12.7	68	12.2	47	23.9	67	20.4	44	42.9	66	27.0	43
128	6.5	132	7.6	75	12.0	134	11.6	77	21.7	131	15.1	77
256	3.4	252	4.1	140	6.2	255	6.6	135	11.4	249	9.0	129
512	1.9	455	2.2	260	3.4	472	3.3	267	6.0	474	4.9	236
1024	1.1	777	1.1	498	1.9	846	1.7	528	3.2	885	2.3	494

Table 3: Strong scaling results: Time (s) and Speedup for $p \in \{2, 3, 4\}$ in two dimensions having approximately 17 Mio. dofs. First row shows results for the cG variant of the IETI-DP method, whereas the second row contains results for the dG version. Each column corresponds to a degree p .

17. G. Guennebaud, B. Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
18. C. Hofer. Analysis of discontinuous Galerkin dual-primal isogeometric tearing and interconnecting methods. Technical Report No. 2016-03, <https://www.dk-compmath.jku.at/publications/dk-reports/2016-11-03>, DK Computational Mathematics Linz Report Series, 2016.
19. C. Hofer and U. Langer. Dual-primal isogeometric tearing and interconnecting methods. In P. Neittanmakki, J. Periaux, and O. Pironneau, editors, *Contributions to PDE for Applications*, Springer-ECCOMAS series “Computational Methods in Applied Sciences”. Springer, Berlin, Heidelberg, New York, 2016. to appear.
20. C. Hofer and U. Langer. Dual-primal isogeometric tearing and interconnecting solvers for multipatch dG-IgA equations. *Computer Methods in Applied Mechanics and Engineering*, 2016. In Press, Accepted Manuscript, <http://dx.doi.org/10.1016/j.cma.2016.03.031>.
21. C. Hofer, U. Langer, and I. Touloupoulos. Discontinuous Galerkin Isogeometric Analysis of Elliptic Diffusion Problems on Segmentations with Gaps. *SIAM Journal on Scientific Computing*, 38(6):A3430–A3460, 2016. available also at: <http://arxiv.org/abs/1511.05715>.
22. C. Hofer, U. Langer, and I. Touloupoulos. Discontinuous Galerkin Isogeometric Analysis on non-matching segmentation: error estimates and efficient solvers. RICAM-Report 23, Johann Radon Institute for Computational and Applied Mathematics, Austrian Academy of Sciences, 2016. available at [https://www.ricam.oeaw.ac.at/publications/ricam-reports/Report No. 2016-23](https://www.ricam.oeaw.ac.at/publications/ricam-reports/Report%20No.%202016-23).
23. C. Hofer and I. Touloupoulos. Discontinuous Galerkin Isogeometric Analysis of elliptic problems on segmentations with non-matching interfaces. *Computers & Mathematics with Applications*, 72(7):1811 – 1827, 2016.
24. J. Hoschek and D. Lasser. *Fundamentals of Computet Aided Geometric Design*. A K Peters, Wellesley, Massachusetts, 1993. Translated by L. Schumaker.
25. T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *Comput. Methods Appl. Mech. Engrg.*, 194:4135–4195, 2005.
26. B. Jüttler, M. Kapl, D.-M. Nguyen, Q. Pan, and M. Pauley. Isogeometric segmentation: The case of contractible solids without non-convex edges. *Computer-Aided Design*, 57:74–90, 2014.
27. A. Klawonn, M. Lanser, and O. Rheinbach. Nonlinear FETI-DP and BDDC methods. *SIAM J. Sci. Comput.*, 36(2):737–765, 2014.
28. A. Klawonn, M. Lanser, and O. Rheinbach. Toward extremely scalable nonlinear domain decomposition methods for elliptic partial differential equations. *SIAM J. Sci. Comput.*, 37(6):c667–c696, 2015.




Fig. 5: Strong scaling of the cG-IETI-DP (left column) and dG-IETI-DP (right column) method for B-Spline degrees $p \in \{2, 3, 4\}$ in three dimensions. The markers $\{\circ, *, \diamond\}$ as well as different shades of red (assembling phase) and blue (solver phase) correspond to the degrees $\{2, 3, 4\}$.

29. A. Klawonn, M. Lanser, and O. Rheinbach. A nonlinear FETI-DP method with an inexact coarse problem. In *Domain decomposition methods in science and engineering XXII. Proceedings of the 22nd international conference on domain decomposition methods, Lugano, Switzerland, September 16–20, 2013*, pages 41–52. Cham: Springer, 2016.
30. A. Klawonn, M. Lanser, O. Rheinbach, H. Stengel, and G. Wellein. *Hybrid MPI/OpenMP Parallelization in FETI-DP Methods*, pages 67–84. Springer International Publishing, Cham, 2015.
31. A. Klawonn and O. Rheinbach. A parallel implementation of Dual-Primal FETI methods for three-dimensional linear elasticity using a transformation of basis. *SIAM Journal on Scientific Computing*, 28(5):1886–1906, 2006.
32. A. Klawonn and O. Rheinbach. Inexact FETI-DP methods. *International journal for numerical methods in engineering*, 69(2):284–307, 2007.
33. A. Klawonn and O. Rheinbach. Highly scalable parallel domain decomposition methods with an application to biomechanics. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 90(1):5–32, 2010.
34. A. Klawonn, O. Rheinbach, and L. F. Pavarino. Exact and inexact FETI-DP methods for spectral elements in two dimensions. In *Domain decomposition methods in science and engineering XVII. Selected papers based on the presentations at the 17th international conference on domain decomposition methods, St. Wolfgang/Strobl, Austria, July 3–7, 2006.*, pages 279–286. Berlin: Springer, 2008.
35. S. Kleiss, C. Pechstein, B. Jüttler, and S. Tomar. IETI-isogeometric tearing and interconnecting. *Computer Methods in Applied Mechanics and Engineering*, 247:201–215, 2012.
36. A. Kuzmin, M. Luisier, and O. Schenk. Fast methods for computing selected elements of the greens function in massively parallel nanoelectronic device simulations. In F. Wolf, B. Mohr, and D. Mey, editors, *Euro-Par 2013 Parallel Processing*, volume 8097 of *Lecture Notes in Computer Science*, pages 533–544. Springer Berlin Heidelberg, 2013.
37. U. Langer, A. Mantzaflaris, S. E. Moore, and I. Touloupoulos. Multipatch discontinuous Galerkin isogeometric analysis. In B. Jüttler and B. Simeon, editors, *Isogeometric Analysis and Applications IGAA 2014*, volume 107 of *Lecture Notes in Computer Science*, pages 1–32, Heidelberg, 2015. Springer. also available at <http://arxiv.org/abs/1411.2478>.
38. U. Langer and I. Touloupoulos. Analysis of multipatch discontinuous Galerkin IgA approximations to elliptic boundary value problems. *Computing and Visualization in Science*, 17(5):217–233, 2015.
39. J. Li and O. B. Widlund. On the use of inexact subdomain solvers for BDDC algorithms. *Comput. Methods Appl. Mech. Eng.*, 196(8):1415–1428, 2007.
40. J. Mandel, C. R. Dohrmann, and R. Tezaur. An algebraic theory for primal and dual substructuring methods by constraints. *Appl. Numer. Math.*, 54(2):167–193, 2005.
41. A. Mantzaflaris, C. Hofer, et al. G+Smo (Geometry plus Simulation modules) v0.8.1. <http://gs.jku.at/gismo>, 2015.
42. M. Pauley, D.-M. Nguyen, D. Mayer, J. Speh, O. Weeger, and B. Jüttler. The isogeometric segmentation pipeline. In B. Jüttler and B. Simeon, editors, *Isogeometric Analysis and Applications IGAA 2014*, Lecture Notes in Computer Science, Heidelberg, 2015. Springer. to appear, also available as Technical Report no. 31 at <http://www.gs.jku.at>.

3d	$p = 2$				$p = 3$				$p = 4$			
cG-IETI-DP	assembling phase		solving phase		assembling phase		solving phase		assembling phase		solving phase	
# procs	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.
8	140.4	8	93.7	8	624.8	8	205.3	8	2565.7	8	382.6	8
16	70.5	16	47.7	16	312.8	16	103.4	16	1285.6	16	182.5	17
32	35.2	32	25.0	30	157.0	32	53.7	31	643.3	32	93.7	33
64	17.5	64	11.9	63	78.6	64	27.2	60	322.5	64	46.6	66
128	9.0	125	7.9	95	40.2	124	14.9	110	163.0	126	26.4	116
256	4.8	236	4.2	178	20.4	245	8.8	187	82.1	250	13.9	221
512	2.5	452	2.6	294	10.3	483	5.6	296	41.4	496	10.0	305
1024	1.4	807	1.5	506	5.5	906	3.2	509	21.4	961	5.8	526
dG-IETI-DP	assembling phase		solving phase		assembling phase		solving phase		assembling phase		solving phase	
# procs	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.	Time	Sp.
8	249.6	8	210.8	8	985.8	8	433.2	8	3588.1	8	854.8	8
16	126.2	16	106.6	16	498.6	16	217.1	16	1792.9	16	405.7	17
32	65.0	31	56.6	30	255.4	31	110.3	31	913.5	31	205.0	33
64	33.1	60	30.5	55	128.5	61	58.2	60	460.0	62	105.9	65
128	17.4	115	17.0	99	65.5	120	30.7	113	234.3	123	56.4	121
256	9.4	212	9.6	175	33.8	233	18.4	188	117.8	244	35.4	193
512	5.1	391	6.1	277	17.4	453	11.5	302	59.9	479	21.4	320
1024	3.1	653	3.5	481	9.6	822	7.1	491	31.3	917	13.2	517

Table 4: Strong scaling results: Time (s) and Speedup for $p \in \{2, 3, 4\}$ in three dimensions having approximately 5 Mio. dofs. First row shows results for the cG variant of the IETI-DP method, whereas the second row contains results for the dG version. Each column corresponds to a degree p .

43. C. Pechstein. *Finite and boundary element tearing and interconnecting solvers for multiscale problems*. Berlin: Springer, 2013.
44. O. Rheinbach. Parallel iterative substructuring in structural mechanics. *Arch. Comput. Methods Eng.*, 16(4):425–463, 2009.
45. B. Rivière. *Discontinuous Galerkin methods for solving elliptic and parabolic equations. Theory and implementation*. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 2008.
46. A. Toselli and O. B. Widlund. *Domain decomposition methods – algorithms and theory*. Berlin: Springer, 2005.
47. X. Tu. Three-level BDDC in three dimensions. *SIAM Journal on Scientific Computing*, 29(4):1759–1780, 2007.
48. X. Tu. Three-level BDDC in two dimensions. *International journal for numerical methods in engineering*, 69(1):33–59, 2007.
49. S. Zampini. *Inexact BDDC Methods for the Cardiac Bidomain Model*, pages 247–255. Springer International Publishing, Cham, 2014.

cG-IETI-DP	$p = 2$			$p = 3$			$p = 4$		
# $S_{\Pi\Pi}^{-1}$ Holder	Assemble Time	Solving Time	Total Time	Assemble Time	Solving Time	Total Time	Assemble Time	Solving Time	Total Time
1	3.61	3.66	7.27	6.50	5.52	12.02	11.23	8.30	19.53
2	4.49	3.58	8.07	7.97	5.57	13.54	13.83	8.02	21.85
4	4.53	3.82	8.35	7.65	5.40	13.05	13.60	8.09	21.69
8	4.46	3.63	8.09	7.72	5.76	13.48	13.32	8.15	21.47
16	4.34	3.49	7.83	7.64	5.61	13.25	13.16	7.93	21.09
32	4.33	3.73	8.06	7.74	5.39	13.13	13.15	8.78	21.93
64	4.34	3.59	7.93	7.62	5.45	13.07	13.10	8.04	21.14
128	4.49	4.06	8.55	7.60	6.05	13.65	13.06	8.47	21.53
256	4.31	4.64	8.95	7.63	6.43	14.06	13.02	8.81	21.83
512	4.34	3.61	7.95	7.55	5.71	13.26	13.23	8.09	21.32
1024	3.73	3.80	7.53	6.56	5.77	12.33	11.19	8.26	19.45

dG-IETI-DP	$p = 2$			$p = 3$			$p = 4$		
# $S_{\Pi\Pi}^{-1}$ Holder	Assemble Time	Solving Time	Total Time	Assemble Time	Solving Time	Total Time	Assemble Time	Solving Time	Total Time
1	4.57	5.09	9.66	7.28	7.16	14.44	12.44	10.01	22.45
2	5.23	4.16	9.39	9.10	6.26	15.36	15.02	9.01	24.03
4	5.25	4.18	9.43	9.12	6.58	15.70	14.93	8.73	23.66
8	5.19	4.28	9.47	8.97	6.29	15.26	14.95	9.30	24.25
16	5.26	4.20	9.46	8.78	6.41	15.19	14.79	9.16	23.95
32	5.11	4.64	9.75	8.82	6.29	15.11	14.96	9.05	24.01
64	5.35	4.75	10.1	9.06	6.87	15.93	14.85	9.37	24.22
128	5.07	6.06	11.13	8.88	8.25	17.13	14.61	10.65	25.26
256	5.07	5.89	10.96	8.66	7.77	16.43	14.52	11.32	25.84
512	5.03	6.15	11.18	8.66	8.29	16.95	14.43	11.16	25.59
1024	4.70	5.33	10.03	7.45	7.68	15.13	12.89	10.60	23.49

Table 5: Influence of the number of processors having an LU-factorization of $S_{\Pi\Pi}$. Timings in seconds for 1024 Processors on a domain with around 70 Mio. dofs and 2048 subdomains.

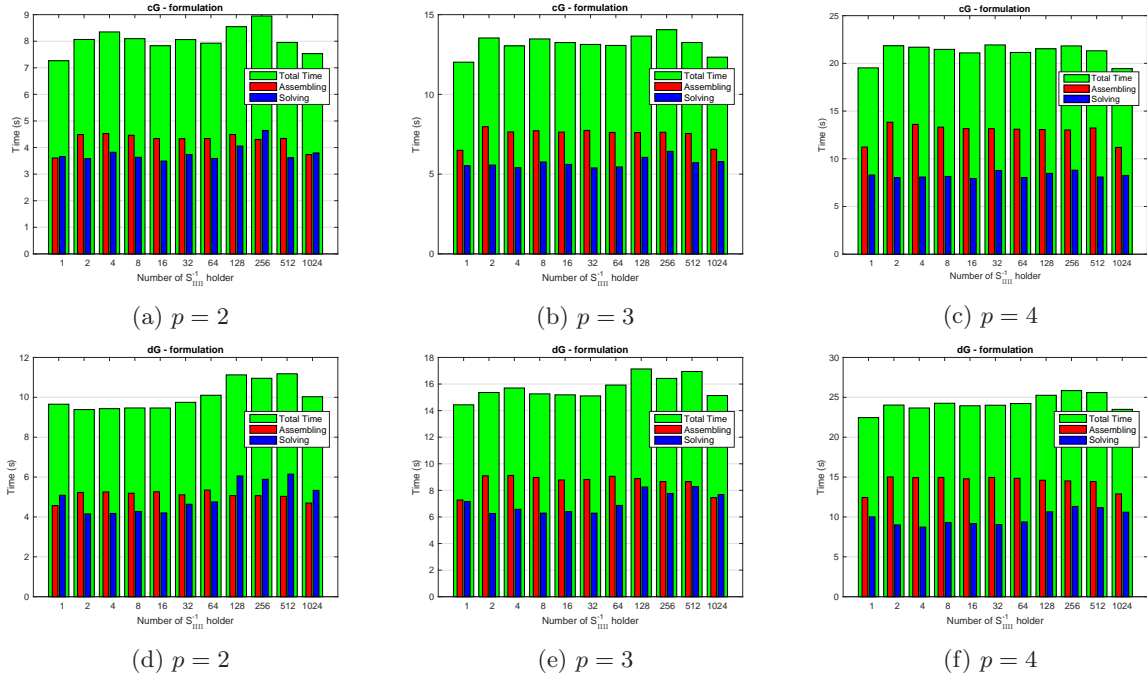


Fig. 6: Influence of the number of $S_{\Pi\Pi}^{-1}$ holders on the scaling. First row corresponds to cG-IETI-DP, second row to dG-IETI-DP. Each column has a fixed degree $p \in \{2, 3, 4\}$. Figures (a-c) summarizes the cG version and Figures (d-f) the dG version, respectively.